



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 079 583 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
28.02.2001 Bulletin 2001/09

(51) Int Cl.: H04L 29/06, H04L 12/24

(21) Application number: 99480077.9

(22) Date of filing: 20.08.1999

(84) Designated Contracting States:
 AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE
 Designated Extension States:
 AL LT LV MK RO SI

(71) Applicant: INTERNATIONAL BUSINESS
 MACHINES CORPORATION
 Armonk, NY 10504 (US)

(72) Inventors:
 • Daude, Olivier
 06200 Nice (FR)
 • Hericourt, Olivier
 06800 Cagnes-sur-Mer (FR)

(74) Representative: Etorre, Yves Nicolas
 Compagnie IBM France
 Département Propriété Intellectuelle
 Le Plan du Bois
 06610 La Gaude (FR)

(54) Method and system for optimizing performance and availability of a dynamic host configuration protocol (DHCP) service

(57) The present invention discloses a method and system for monitoring and optimizing performance and availability of a Dynamic Host Configuration Protocol (DHCP) service provided by one or a plurality of DHCP servers (602) in an Internet Protocol (IP) network comprising one or a plurality of IP subnetworks.

The method comprises in a DHCP server (602) the steps of:

- defining one or a plurality of groups of subnetworks, a group of subnetworks comprising one or a plurality of subnetworks;
- retrieving information related to resources, in particular IP addresses, allocated within said DHCP server to each group of subnetworks;
- transferring said information to a DHCP service monitoring system (600).

The method comprises in a DHCP service monitoring system (403) the steps of:

- retrieving (501 to 505) from one or a plurality of DHCP servers (401), information related to resources, in particular IP addresses, allocated within each DHCP server (401) to groups of subnetworks, each group of subnetworks comprising one or a plurality of subnetworks;
- aggregating (506 to 511) the information for each group of subnetworks.

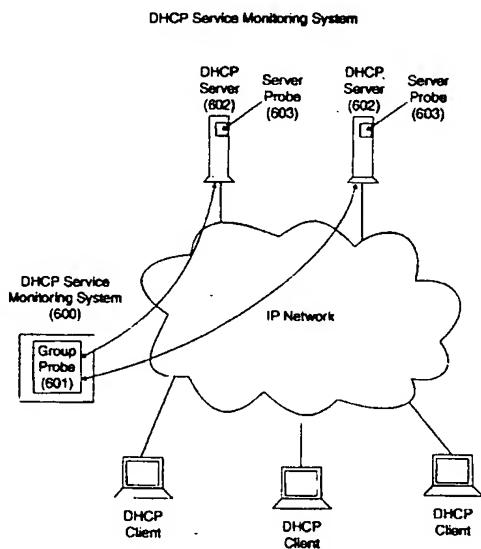


FIG. 5

Description**Technical field of the invention**

[0001] The present invention relates to computer networks, and more particularly to a method and system in an Internet Protocol (IP) network for optimizing performance and availability of a Dynamic Host Configuration Protocol (DHCP) service provided by a plurality of DHCP servers.

Background art**INTERNET**

[0002] Internet is a global network of computers and computers networks (the "Net"). The Internet connects computers that use a variety of different operating systems or languages, including UNIX, DOS, Windows, Macintosh, and others. To facilitate and allow the communication among these various systems and languages, the Internet uses a language referred to as TCP/IP ("Transmission Control Protocol/Internet Protocol"). TCP/IP protocol supports three basic applications on the Internet:

- transmitting and receiving electronic mail,
- logging into remote computers (the "Telnet"), and
- transferring files and programs from one computer to another ("FTP" or "File Transfer Protocol").

[0003] One of the object of TCP/IP is to interconnect networks and to provide an universal communication services: an inter-network, or Internet. Each physical network has its own technology-dependent communication interface, in the form of a programming interface that provides basic communication functions (primitives). Communication services are provided by a software that runs between the physical network and user applications. This software provides a common interface for these applications, independent of the underlying physical network. The architecture of the physical networks is hidden from the user.

[0004] The Internet protocol is still evolving through the mechanism of Request For Comments (RFC). New protocols (mostly application protocols) are designed and implemented by researchers. They are brought to the attention of the Internet community in the form of an Internet Draft (ID). The largest source of IDs is the Internet Engineering Task Force (IETF).

IP ADDRESSES

[0005] To interconnect two networks, a computer system able to forward packets from one network to the other is attached to both networks. Such a machine is called a router. The term "IP router" is also used because the routing function is part of the IP layer of the TCP/IP pro-

tocol.

[0006] IP addresses are used by the IP protocol to uniquely identify a host on the Internet (Strictly speaking, an IP address identifies an interface that is capable of sending and receiving IP datagrams, and one system can have multiple such interfaces. However, both hosts and routers must have at least one IP address, so this simplified definition is acceptable). IP datagrams (the basic data packets exchanged between hosts) are transmitted by a physical network attached to the host and each IP datagram contains a source IP address and a destination IP address.

[0007] IP addresses are represented by a 32-bit unsigned binary value which is usually expressed in a dotted decimal format. For example, 9.167.5.8 is a valid Internet address. The numeric form is used by the IP software. The mapping between the IP address and an easier-to-read symbolic name, for example myhost.ibm.com, is done by the Domain Name System

IP SUBNETS

[0008] Due to the explosive growth of the Internet, the principle of assigned IP addresses became too inflexible to allow easy changes to local network configurations. Those changes might occur when:

- A new type of physical network is installed at a location.
- Growth of the number of hosts requires splitting the local network into two or more separate networks.
- Growing distances require splitting a network into smaller networks, with gateways between them.

[0009] To avoid having to request additional IP network addresses in these cases, the concept of subnets was introduced. The assignment of subnets can be done locally, as the whole network still appears to be one IP network to the outside world.

[0010] The host number part of the IP address is subdivided again into a network number and a host number. This second network is termed a subnetwork or subnet. The main network now consists of a number of subnets and the IP address is interpreted as:

<network number><subnet number><host number>

[0011] The combination of the subnet number and the host number is often termed the local address or the local part. Subnetting is implemented in a way that is transparent to remote networks. A host within a network that has subnets is aware of the subnetting but a host in a different network is not; it still regards the local part of the IP address as a host number.

[0012] The division of the local part of the IP address into subnet number and host number parts can be chosen freely by the local administrator; any bits in the local part can be used to form the subnet. The division is done using a subnet mask which is a 32 bit number. Zero bits

in the subnet mask indicate bit positions ascribed to the host number, and ones indicate bit positions ascribed to the subnet number. The bit positions in the subnet mask belonging to the network number are set to ones but are not used. Subnet masks are usually written in dotted decimal form, like IP addresses.

DOMAIN NAMES

[0013] The host or computer names (like www.enterprise.com) are translated into numeric Internet addresses (like 194.56.78.3), and vice versa, by using a method called DNS ("Domain Name Service"). DNS is supported by network-resident servers, also known as domain name servers or DNS servers.

DYNAMIC IP

[0014] There are generally three pieces of information needed by a system in order to be able to communicate on a TCP/IP network:

- an IP address (to uniquely identify the system on the network),
- a subnet mask (to determine the network and subnet parts of the address), and
- the address of at least one router (if the system is to be able to communicate with other devices outside of its immediate subnet).

[0015] These three values represent the bare minimum of information that must be programmed into each and every device for participating in the TCP/IP world. Often the number of necessary parameters will be much higher. With the exponential growth rate of networking today, it is easy to see that manual programming of these values into every device that is to attach to the network represents a major administrative workload.

[0016] The increasingly mobile nature of the end users also raises problems with regard to configuration of network devices. It is possible to allocate multiple sets of configuration parameters to a device, but:

- this obviously means even more workload for the administrator,
- this is wasteful with respect to the number of IP addresses allocated.

[0017] Several components of TCP/IP help automate device configuration, reduce the number of IP addresses allocated, and/or cope with the demands of the mobile user.

BOOTSTRAP PROTOCOL (BOOTP)

[0018] The BOOTP protocol was originally developed as a mechanism to enable diskless hosts to be remotely booted over a network as workstations, routers, terminal

concentrators and so on. It allows a minimum IP protocol stack with no configuration information to obtain enough information to begin the process of downloading the necessary boot code. BOOTP does not define how the downloading is done, but this process typically uses TFTP "Trivial File Transfer Protocol (TFTP)" as described in RFC 906 - Bootstrap Loading Using TFTP. Although still widely used for this purpose by diskless hosts, BOOTP is also commonly used solely as a mechanism to deliver configuration information to a client that has not been manually configured. The BOOTP process involves the following steps:

- 1. The client determines its own hardware address; this is normally in a ROM (Read Only Memory) on the hardware.
- 2. A BOOTP client sends its hardware address in a UDP (User Datagram Protocol) datagram to the server. If the client knows its IP address and/or the address of the server, it should use them, but in general BOOTP clients have no IP configuration data at all. If the client does not know its own IP address, it uses 0.0.0.0. If the client does not know the server's IP address, it uses the limited broadcast address (255.255.255.255). The UDP port number is 67.
- 3. The server receives the datagram and looks up the hardware address of the client in its configuration file, which contains the client's IP address. The server fills in the remaining fields in the UDP datagram and returns it to the client using UDP port 68.
- 4. When it receives the reply, the BOOTP client will record its own IP address and begin the bootstrap process.

[0019] BOOTP is a draft standard protocol. Its status is recommended. The BOOTP specifications can be found in RFC 951 - Bootstrap Protocol. There are also updates to BOOTP, some relating to inter operability with DHCP (Dynamic Host Configuration Protocol), described in RFC 1542 - Clarifications and Extensions for the Bootstrap Protocol, which updates RFC 951 and RFC 2132 - DHCP Options and BOOTP Vendor Extensions.

DYNAMIC HOST CONFIGURATION PROTOCOL (DHCP)

[0020] The Dynamic Host Configuration Protocol (DHCP) provides a framework for passing configuration information to hosts on a TCP/IP network. DHCP is based on the BOOTP protocol, adding the capability of automatic allocation of reusable network addresses and additional configuration options. DHCP messages use UDP port 67, the BOOTP server's well-known port and

UDP port 68, the BOOTP client's well-known port. DHCP consists of two components:

- 1. A protocol that delivers host-specific configuration parameters from a DHCP Server to a host.
- 2. A mechanism for the allocation of temporary or permanent network addresses to hosts.

[0021] IP requires the setting of many parameters within the protocol implementation software. Because IP can be used on many dissimilar kinds of network hardware, values for those parameters cannot be guessed at or assumed to have correct default values. The use of a distributed address allocation scheme based on a polling / defence mechanism, for discovery of network addresses already in use, cannot guarantee unique network addresses because hosts may not always be able to defend their network addresses. DHCP supports three mechanisms for IP address allocation:

- 1. Automatic allocation : DHCP assigns a permanent IP address to the host.
- 2. Dynamic allocation : DHCP assigns an IP address for a limited period of time.
- 3. Manual allocation : The host's address is assigned by a network administrator.

[0022] DHCP is a draft standard protocol. Its status is elective. The current DHCP specifications can be found in RFC 2131 - Dynamic Host Configuration Protocol and RFC 2132 - DHCP Options and BOOTP Vendor Extensions.

Configuration Parameters Repository

[0023] DHCP provides persistent storage of network parameters for network clients. A DHCP Server stores a key-value entry for each client, the key being some unique identifier, for example an IP subnet number and a unique identifier within the subnet (normally a hardware address), and the value contains the configuration parameters last allocated to this particular client.

[0024] One effect of this is that a DHCP client will tend to always be allocated the same IP address by the server, provided the pool of addresses is not over-subscribed and the previous address has not already been allocated to another client.

DHCP Considerations

[0025] DHCP dynamic allocation of IP addresses and configuration parameters relieves the network administrator of great deal of manual configuration work. The ability for a device to be moved from network to network and to automatically obtain valid configuration parameters for the current network can be of great benefit to mobile users. Also, because IP addresses are only allocated when clients are actually active, it is possible,

by the use of reasonably short lease times and the fact that mobile clients do not need to be allocated more than one address, to reduce the total number of addresses in use in an organization. However, the following should be considered when DHCP is being implemented:

- DHCP is built on UDP, which is, as yet, inherently insecure. In normal operation, an unauthorized client could connect to a network and obtain a valid IP address and configuration. To prevent this, it is possible to preallocate IP addresses to particular MAC (Medium Access Control) addresses (similar to BOOTP), but this increases the administration workload and removes the benefit of recycling of addresses.
- Unauthorized DHCP Servers could also be set up, sending false and potentially disruptive information to clients.
- In a DHCP environment where automatic or dynamic address allocation is used, it is generally not possible to predetermine the IP address of a client at any particular point in time. In this case, if static DNS (Domain Name Server) servers are also used, the DNS servers will not likely contain valid host name to IP address mappings for the clients. If having client entries in the DNS is important for the network, one may use DHCP to manually assign IP addresses to those clients and then administer the client mappings in the DNS accordingly.

BOOTP and DHCP Inter operability

[0026] The format of DHCP messages is based on the format of BOOTP messages, which enables BOOTP and DHCP clients to interoperate in certain circumstances. Support for BOOTP clients at a DHCP Server must be configured by a system administrator, if required.

DYNAMIC DOMAIN NAME SYSTEM

[0027] In order to take advantage of DHCP, yet still to be able to locate any specific host by means of a meaningful label, such as its host name, the following extensions to the Domain Name System (DNS) are required:

- A method for the host name to address mapping entry for a client in the domain name server to be updated, once the client has obtained an address from a DHCP Server.
- A method for the reverse address to host name mapping to take place once the client obtains its address.
- Updates to the DNS to take effect immediately, without the need for intervention by an administrator.

- Updates to the DNS to be authenticated to prevent unauthorized hosts from accessing the network and to stop imposters from using an existing host name and remapping the address entry for the unsuspecting host to that of its own.
- A method for primary and secondary DNS servers to quickly forward and receive changes as entries are being updated dynamically by clients. In short, a secure Dynamic Domain Name System (DDNS) is necessary.

[0028] In summary, in the DHCP and DDNS environment, DHCP provides a device with a valid IP address for the point at which it is attached to the network. DDNS provides a method of locating that device by its host name, no matter where that device happens to be attached to a network and what IP address it has been allocated.

[0029] More explanations about the domain presented in the above sections can be found in the following publications incorporated herewith by reference:

- TCP/IP Tutorial and Technical Overview by Martin W. Murhammer, Orcun Atakan, Stefan Bretz, Larry R. Pugh, Kazunari Suzuki, David H. Wood published by IBM International Technical Support Organization.
- "Internet in a nutshell" by Valerie Quercia, published by O'Reilly, October 1997.

PROBLEM

[0030] The problem is to monitor the utilization of a Dynamic Host Configuration Protocol (DHCP) service provided by one or a plurality of DHCP servers, in order to optimally adjust configuration parameters. Due to the nature of DHCP protocol which is based on UDP BOOTP broadcast, the DHCP service is provided to the DHCP Client by the "fastest" DHCP Server to answer. As a consequence, in a multiple DHCP Servers environment, when a DHCP Server fails or when a DHCP Server reaches its maximum capacity, the DHCP Servers which are still available continue to provide a DHCP Service and to answer DHCP Client requests. The available DHCP Servers continue to provide the service until they reach their maximum capacity. Eventually the DHCP Service may fail because available DHCP Servers cannot support the total load.

[0031] A single DHCP Server may be configured to provide a DHCP Service to multiple subnetworks or groups of subnetworks. In this case, the DHCP Server is configured with one pool of IP addresses per subnetwork. The capacity or size of a pool is defined by the number of IP addresses this pool can manage.

[0032] In addition, for better resilience and better performance, a group of subnetworks can be administered by multiple DHCP Servers. In this case, the DHCP Serv-

ice is provided by this plurality of DHCP Servers.

[0033] Note : in the present invention, "a group of subnetworks" comprises one or multiple subnetworks also called "IP subnets" generally located in a same geographical area.

[0034] The problems are then to:

- monitor the utilization of the DHCP Service for each group of subnetworks.
- monitor the pools of IP addresses configured within each DHCP Server.
- correlate and monitor for each group of subnetworks, the utilization of the pools of IP addresses configured on multiple DHCP Servers
- detect any capacity problem on DHCP Servers, in order to anticipate any DHCP Service degradation or disruption.

[0035] The DHCP Service provided when one DHCP Server has reached its maximum capacity is degraded because the DHCP Servers which are still available have to handle additional load. The monitoring of a DHCP Service without analysis of its utilisation cannot not lead to an efficient anticipation of the problems, in particular disruption or degradation of the DHCP Service.

Objects of the invention

[0036]

- One object of the present invention is to provide a method for optimizing performance and availability of a DHCP Service.
- It is a further object of the present invention to monitor the utilisation of a DHCP Service for each group of subnetworks.
- It is another object of the present invention to monitor the pools of IP addresses configured within each DHCP Server.
- It is yet another object of the present invention to correlate and monitor the utilization of the pools of IP addresses configured on multiple DHCP Servers for each group of subnetworks.
- It is yet another object of the present invention to detect any capacity problem on DHCP Servers.
- It is another object of the present invention to check that every DHCP Server is correctly configured.

Summary of the invention

[0037] The present invention discloses a method and system for monitoring and optimizing performance and availability of a Dynamic Host Configuration Protocol

(DHCP) service provided by one or a plurality of DHCP servers in an Internet Protocol (IP) network comprising one or a plurality of IP subnetworks.

[0038] More particularly the present invention uses specific probes for collecting information related to the utilization of IP address pools within said DHCP Servers.

[0039] The present invention also uses a DHCP Service Utilisation application for correlating the information related to the utilization of the IP address pools, and providing monitoring reports and statistics.

[0040] The method comprises in a DHCP server the steps of:

- defining one or a plurality of groups of subnetworks, each group of subnetworks comprising one or a plurality of subnetworks;
- retrieving information related to resources, in particular IP addresses, allocated within said DHCP server to each group of subnetworks;
- transferring said information to a DHCP service monitoring system.

[0041] The method comprises in a DHCP service monitoring system the steps of:

- retrieving from one or a plurality of DHCP servers, information related to resources, in particular IP addresses, allocated within each DHCP server to groups of subnetworks, each group of subnetworks comprising one or a plurality of subnetworks;
- aggregating the information for each group of subnetworks.

[0042] The present inventions provides the following advantages:

- Early detection of IP address shortage in pools configured within each DHCP Server for each group of subnetworks.
- Early detection of IP address shortage in pools configured on multiple DHCP Servers for each group of subnetworks.
- Early detection of any capacity problem on DHCP Servers for each group of subnetworks in order to anticipate any service degradation or disruption.
- Production of DHCP Service utilisation reports, for adjusting configuration of DHCP Servers (for example, increase of the size of an IP address pool for a specific group of subnetworks on a DHCP Server).
- No additional or specific hardware is mandatory.

Drawings

[0043] The novel and inventive features believed characteristics of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objects and advantages

thereof, will best be understood by reference to the following detailed description of an illustrative detailed embodiment when read in conjunction with the accompanying drawings, wherein:

- 5 • Figure 1 shows the IP address allocation to a DHCP Client by a DHCP Server according to prior art.
- 10 • Figure 2 is a view of a DHCP server probe according to the present invention.
- Figure 3 is a flow chart of a DHCP server probe according to the present invention.
- 15 • Figure 4 is a view of a DHCP group probe according to the present invention.
- Figure 5 is a flow chart of a DHCP group probe according to the present invention.
- 20 • Figure 6 is a physical view of a DHCP service monitoring system according to the present invention.

Preferred embodiment of the invention

[0044] The present invention relates to the monitoring of a DHCP service and more particularly to a method and system for checking the size of IP address pools for a group of IP subnetworks among one or a plurality of 30 DHCP servers. The present invention is based on the measure of utilization of IP address pools by means of probes.

IP ADDRESS ALLOCATION

[0045] Figure 1 describes the DHCP Client/Server interactions when the DHCP Client does not know its network address. More particularly, Figure 1 shows the acquisition mechanism by a DHCP Client of the IP address and the IP minimal configuration parameters from a DHCP Server within an IP network. The DHCP Client (101) broadcasts a request on its local physical subnet (103). The request is forwarded by any router having a BOOTP forwarding mechanism. When the request is received by a DHCP Server (102), the DHCP Server checks whether it is able to answer the DHCP Client or not. If the DHCP Server has still some available IP address within its address database, a positive answer is returned to the DHCP Client. The DHCP Client selects the first DHCP Server for which a positive answer is received and confirms to this server its agreement.

[0046] More particularly the allocation of a new network address comprises the following steps:

- 55 • The DHCP Client broadcasts a request on its local physical subnet. The request may include some options such as network address suggestion or lease duration.

- Each DHCP Server may respond with a message that includes an available network address and other configuration options. The DHCP Server may record the address as offered to the DHCP Client to prevent the same address being offered to other DHCP Clients in the event of further messages being received before the first DHCP Client has completed its configuration.
- The DHCP Client receives one or more messages from one or more DHCP Servers. The DHCP Client chooses one based on the configuration parameters offered and broadcasts a message that includes the DHCP Server identifier option to indicate which message it has selected and the requested IP address option, taken from the DHCP Client IP address in the selected offer.
- The DHCP Servers receive the messages broadcasted by the DHCP Client. Those DHCP Servers not selected use the message as notification that the DHCP Client has declined that DHCP Server's offer. The DHCP Server selected in the message commits the binding for the DHCP Client to persistent storage and responds with a message containing the configuration parameters for the requesting DHCP Client.
- The DHCP Client receives the message with configuration parameters and performs a final check on the parameters. At this point the DHCP Client is configured.

INTERNAL VIEW OF A DHCP SERVER PROBE

[0047] Figure 2 is a logical view of a DHCP Server Probe (202). The DHCP Server Probe runs concurrently with the DHCP Server application (201). The DHCP Server Probe (202) within the DHCP Server (200) determines, the current number of allocated IP addresses for each group of subnetworks and saves this information in a file (203) called "Server Utilisation File" (in a preferred embodiment, there is one entry in the Server Utilisation File for each group of subnetworks). A group of subnetworks is in fact a list of IP subnetworks (or IP subnets) stored in a file (204) generally called "Group File" (in a preferred embodiment, the file name is the group name for a better identification of the information). A workstation broadcasting an IP address allocation request from its local subnetwork, will be served by a unique DHCP Server. The Group File (204) in a DHCP Server comprises in fact the list of IP subnetworks belonging to a same group.

DYNAMIC VIEW OF A DHCP SERVER PROBE

[0048] Figure 3 is a flow chart of a DHCP Server Probe as described in Figure 2. The DHCP Server

Probe is executed at given and regular period of times. Information collected by the probe is recorded in a file (Server Utilisation File - 203) containing for each group of subnetworks, the total number of IP addresses within the pool and the number of IP addresses used or allocated. The process comprises the following steps:

- (301) The Server Utilization File in the DHCP Server is opened.
- (302) A test determines whether there is a Group File to process or not.
- (313) If there is no more Group File to process, the Server Utilization File in the DHCP Server is closed.
- If there is still a Group File to process:
 - (303) The group name is extracted from the Group File (in a preferred embodiment the group name is extrapolated from the Group File name).
 - (304) The list of subnetworks that belongs to the group is extracted from the Group File.
 - (305) A first counter called "Total Group Pool Size counter" is initialized for the group. This counter is used to determine the total number of IP addresses that can be allocated to this group (available and allocated IP addresses).
 - (306) A second counter called "Total Group Pool Used counter" is initialized for the group. This counter is used to determine the total number of IP addresses actually allocated to this group.
 - (307) A test determines whether there is a subnetwork within the group to process or not.
 - If there is a subnetwork to process:
 - (308) The size of the pool associated with the subnetwork is retrieved from the DHCP Server Application.
 - (309) The size of the pool is added to the Total Group Pool Size stored in the Total Group Pool Size counter.
 - (310) The number of allocated IP addresses associated with the subnetwork is retrieved from the DHCP Server Application.
 - (311) The number of allocated IP addresses is added to the Total Group Pool Used stored in the Total Group Pool Used counter.

- When all subnetworks have been processed:
 - (312) The Server Utilisation File is updated with the following information:
 - name of the group;
 - Total Group Pool Used;
 - Total Group Pool Size;
 - a time stamp.
- If there is a DHCP Server to process:
 - (503) A connection is established with the DHCP Server (401).
 - (504) The Server Utilisation File (402) is retrieved from the DHCP Server.
 - (505) The connection with the DHCP Server is terminated.
- When there is no more DHCP Server to process:
 - (506) An exhaustive list of the groups is extrapolated from the Server Utilisation Files previously retrieved from DHCP Servers (note each Server Utilisation File contains group names, and a same group name may be present in multiple Server Utilisation Files).
 - (507) A test determines whether there is still a group to process or not:
 - If there is still a group to process:
 - (508) Data for the group (Total Group Pool Used, Total Group Pool Size) are retrieved from Server Utilisation Files.
 - (509) The global utilization of the IP address pools for the group is computed. In a preferred embodiment, the global utilisation is equal to sum of Total Group Pool Used variables, divided by sum of all Total Group Pool Size variables.
 - An action (for instance an alert) can be sent if a predefined threshold has been reached for this group.
 - (510) The global utilisation of the IP address pools for the group is recorded with a time stamp.
 - (511) When there is no more group to process, a successful return code is returned.

DHCP GROUP PROBE

[0049] Figure 4 is a view of a DHCP Group Probe (405). The probe can be processed either in a DHCP Monitoring System (403) or in any DHCP Server. The purpose of a DHCP Group Probe is to compute for each group of subnetworks administered by one or a plurality of DHCP Servers, the global utilisation of the IP address pools. The DHCP Group Probe collects the information previously processed by DHCP Server Probes on DHCP Servers (401).

- The information is stored in the Server Utilization File (402) of each DHCP Server (401) administering the group.
- DHCP Servers are recorded in a list called "DHCP Servers List File" (404). The collected information is then aggregated in several files (one per group) called "Group Utilisation Files" (406). Each file contains all information related to the utilization of IP address pools for a given group, in particular the percentage of utilisation of IP addresses. To avoid any IP address shortage in IP address pools, specific actions (for instance, increase of a pool size in a DHCP Server) can then be launched if, for instance, the computed percentage exceeds a pre-defined threshold.

DYNAMIC VIEW OF A DHCP GROUP PROBE

[0050] Figure 5 is a flow chart of a DHCP Group Probe. The probe is preferably executed at given and regular periods of time. The information for each group of subnetworks is collected in the different DHCP Servers. A file (404) in the DHCP Monitoring System (403) contains an exhaustive list of the DHCP Servers administering each group of subnetworks. The following process in the DHCP Monitoring System (403) is then executed:

- (501) A DHCP Server (in a preferred embodiment, the DHCP Server name, or address) is retrieved from the DHCP Servers List File (404).
- (502) A test determines whether there is still a DHCP Server to process or not.

40

45

50

55

55

DHCP SERVICE MONITORING SYSTEM

[0051] Figure 6 is a view of a DHCP Service Monitoring system (600) within an IP network. A dedicated computer system can be used to run the DHCP Group Probes (601) and to collect the information located in each DHCP server (602) by DHCP Server Probes (603).

[0052] While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood that various changes in form and detail may be made therein without departing

from the spirit, and scope of the invention.

Claims

1. A method for monitoring and optimizing performance and availability of a Dynamic Host Configuration Protocol (DHCP) service provided by one or a plurality of DHCP servers (602) in an Internet Protocol (IP) network comprising one or a plurality of IP subnetworks, said method comprising in a DHCP server (602) the steps of:
 - defining one or a plurality of groups of subnetworks, each group of subnetworks comprising one or a plurality of subnetworks;
 - retrieving information related to resources, in particular IP addresses, allocated within said DHCP server to each group of subnetworks;
 - transferring said information to a DHCP service monitoring system (600).
2. The method according to the preceding claim wherein the step of retrieving information related to resources allocated to each group of subnetworks comprises the further step of:
 - determining a total number of IP addresses presently allocated to each group of subnetworks, the total number of IP addresses of a group of subnetworks being equal to the sum of IP addresses presently allocated to each subnetwork defined within said group of subnetworks.
3. The method according to any one of the preceding claims wherein the step of retrieving information related to resources allocated to each group of subnetworks comprises the further step of:
 - determining a total number of IP addresses that can be potentially allocated to each group of subnetworks, the total number of IP addresses of a group of subnetworks being equal to the sum of IP addresses that can be potentially allocated to each subnetwork defined within said group of subnetworks.
4. The method according to any one of the preceding claims comprising the further step of:
 - storing the total number of IP addresses presently allocated to each group and/or that can be potentially allocated to each group of subnetworks in a file (203) within the DHCP server.
5. A system, in particular a DHCP server (200), adapted for carrying out the method according to any one of the preceding claims.
6. A computer readable medium comprising instructions adapted for carrying out the method according to any one of steps 1 to 4.
7. A method for monitoring and optimizing performance and availability of a Dynamic Host Configuration Protocol (DHCP) service provided by one or a plurality of DHCP servers (401) according to claim 5 in an Internet Protocol (IP) network comprising one or a plurality of IP subnetworks, said method comprising in a DHCP service monitoring system (403) the steps of:
 - retrieving (501 to 505) from one or a plurality of DHCP servers (401), information related to resources, in particular IP addresses, allocated within each DHCP server (401) to groups of subnetworks, each group of subnetworks comprising one or a plurality of subnetworks;
 - aggregating (506 to 511) said information for each group of subnetworks.
8. The method according to the preceding claim wherein the retrieved information comprises for each DHCP server a total number of IP addresses presently allocated to each group of subnetworks, the total number of IP addresses of a group of subnetworks in a DHCP server being equal to the sum of IP addresses presently allocated to each subnetwork defined within said group of subnetworks.
9. The method according to any one of claims 7 to 8 wherein the retrieved information comprises for each DHCP server a total number of IP addresses that can be potentially allocated to each group of subnetworks, the total number of IP addresses of a group of subnetworks in a DHCP server being equal to the sum of IP addresses that can be potentially allocated to each subnetwork defined within said group of subnetworks.
10. The method according to any one of claims 7 to 9 comprising a first step of:
 - recording said one or plurality of DHCP servers in a list (404);
11. The method according to any one of claims 7 to 10 wherein the step of aggregating (506 to 511) said information for each group of subnetworks comprises the further step of:
 - building (506) a list of groups of subnetworks using information retrieved from DHCP servers.

12. The method according to any one of claims 7 to 11
wherein the step of aggregating (506 to 511) said
information for a group of subnetworks comprises
the further step of:

- determining a total number of IP addresses
presently allocated to said group of subnet-
works on all DHCP servers, the total number of
IP addresses being equal to the sum of IP ad-
dresses presently allocated to said group of
subnetworks in each DHCP server.

13. The method according to any one of claims 7 to 12
wherein the step of aggregating (506 to 511) said
information for a group of subnetworks comprises
the further step of:

- determining a total number of IP addresses that
can be potentially allocated to said group of
subnetworks on all DHCP servers, the total
number of IP addresses being equal to the sum
of IP addresses that can be potentially allocat-
ed to said group of subnetworks in each DHCP
server.

14. The method according to any one of claims 7 to 13
wherein the step of aggregating (506 to 511) said
information for a group of subnetworks comprises
the further step of:

- dividing the total number of IP addresses pres-
ently allocated to said group of subnetworks on
all DHCP servers by the total number of IP ad-
dresses that can be potentially allocated to said
group of subnetworks on all DHCP servers;
- comparing the result of the division with a pre-
determined threshold;
- triggering a corrective action when this thresh-
old is reached for a group of subnetworks.

15. The method according to any one of claims 7 to 14
wherein the step of triggering a corrective action
comprises the further step of:

- triggering an alert; or/and
- adjusting within one or a plurality of DHCP serv-
ers the number of IP addresses that can poten-
tially be allocated to said group.

16. The method according to any one of claims 7 to 15
wherein the step of retrieving information is execut-
ed at predefined or/and regular periods of time on
request of the DHCP service monitoring system
(403).

17. A system, in particular a DHCP service monitoring
system (403) adapted for carrying out the method
according to any one of claims 7 to 16.

5

10

15

20

25

30

35

40

45

55

IP Address Allocation

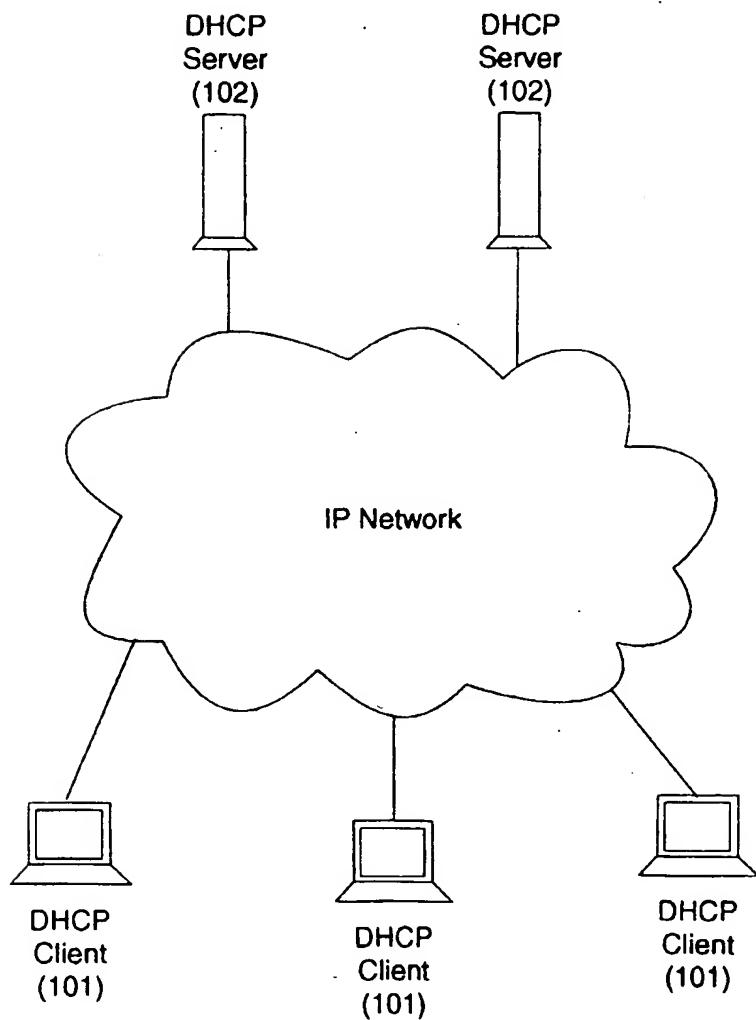


FIG. 1

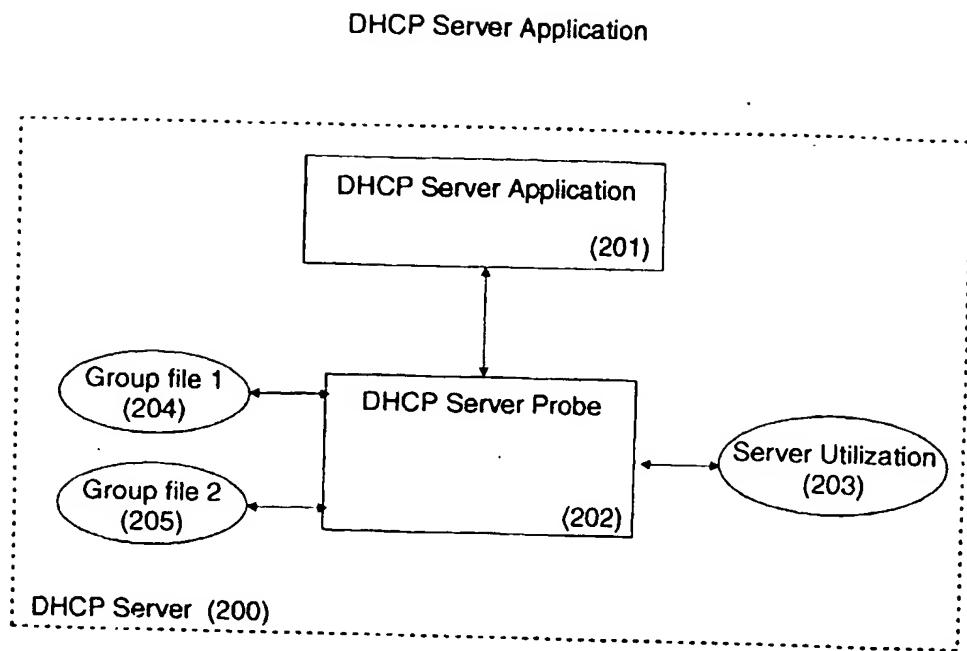


FIG. 2

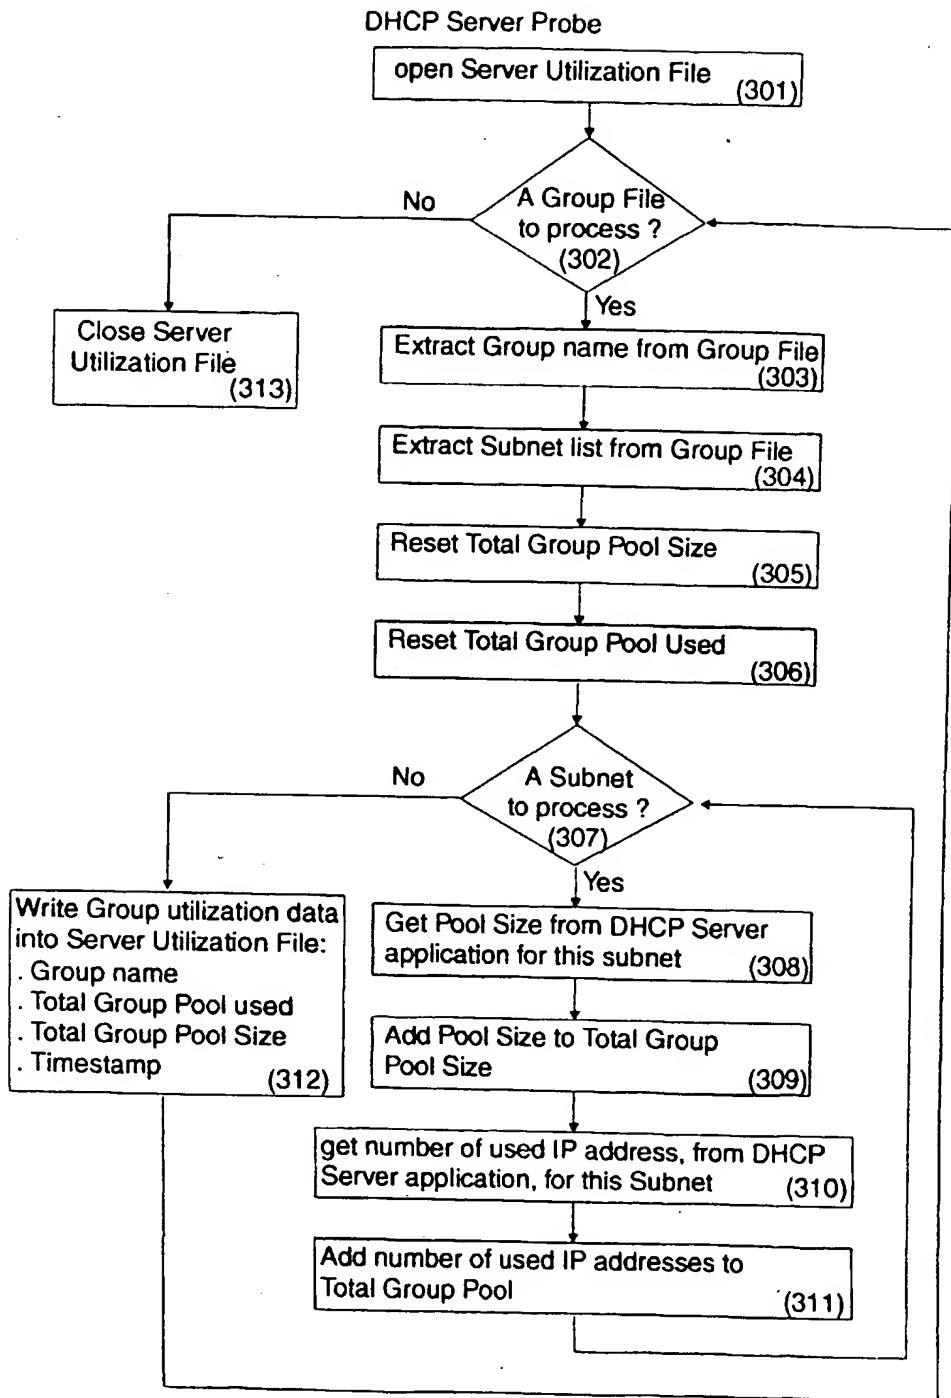


FIG. 3

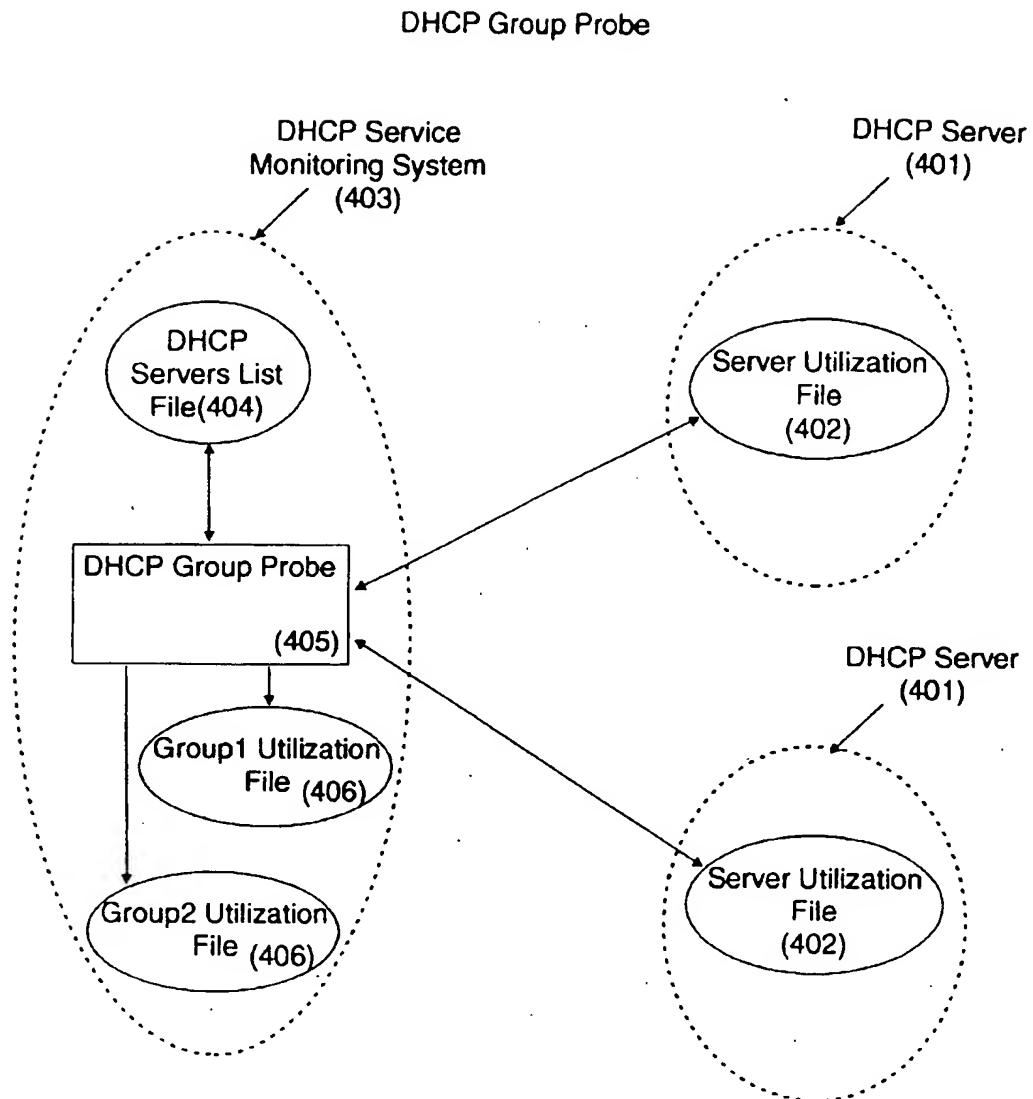


FIG. 4

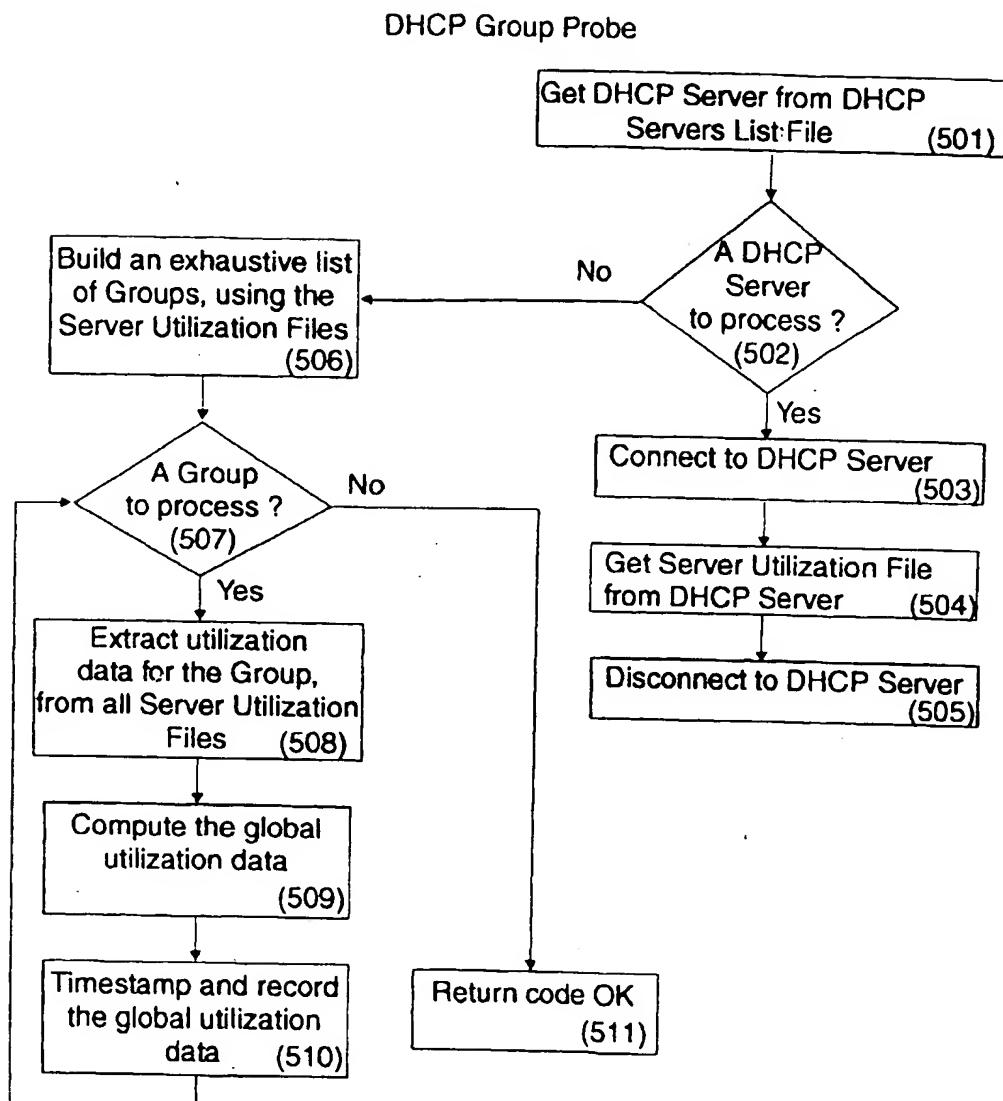


FIG. 5

DHCP Service Monitoring System

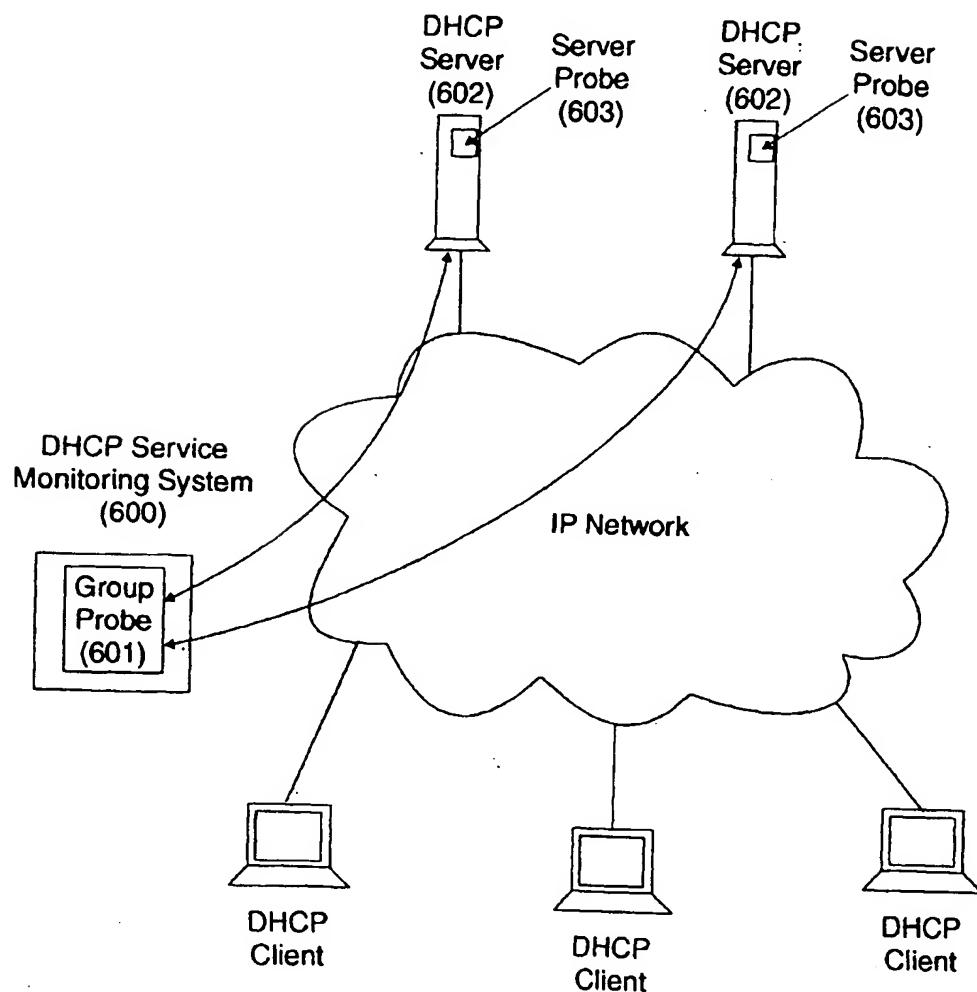


FIG. 6



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 48 0077

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (InCL7)
X	"IP AddressWorks" ONSIGTH SOLUTIONS, 'Online! 1998, XP002129978 Retrieved from the Internet: <URL: http://www.ipworks.nl/ipaw-spd.htm> 'retrieved on 2000-02-07! A * the whole document *	1-13, 17, 18	H04L29/06 H04L12/24
A	WO 99 33211 A (U S WEST INC ;MEDIAONE GROUP INC (US)) 1 July 1999 (1999-07-01) * page 5, line 10 - page 14, line 15 *	14-16	
A	DROMS R: "Automated configuration of TCP/IP with DHCP" IEEE INTERNET COMPUTING, JULY-AUG. 1999, IEEE, USA, vol. 3, no. 4, pages 45-53, XP000874503 ISSN: 1089-7801 * the whole document *	1-18	
TECHNICAL FIELDS SEARCHED (InCL7)			
H04L			
The present search report has been drawn up for all claims			
Place of search THE HAGUE	Date of completion of the search 8 February 2000	Examiner RAMIREZ DE AREL., F	
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : Intermediate document			
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.

EP 99 48 0077

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EPO file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

08-02-2000

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9933211 A	01-07-1999	US 6009103 A	28-12-1999

EPO FORM P048

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82